

## Keynote Presentation

### **Building a Scalable Music Recommender**

**Douglas Eck**, *Research Scientist, Google*

I will discuss techniques for content-based music recommendation at large scale. I'll focus on three steps in a music recommendation pipeline. First I'll present Stabilized Auditory Images (SAIs), acoustic representations based on human auditory processing. Next, I'll look at generating sparse, noise-resistant codes from SAIs using an ordinal coding method called Winner-Take-All (WTA). Finally, I'll describe a multi-class multi-label embedding algorithm, Wsabie for ranking tracks vis-a-vis training labels such as genre and style. The work I discuss is based on work published by a number of researchers at Google. Time permitting, I hope to present several audio demos to illustrate our approach.

Before joining Google in 2011, Douglas Eck was an Associate Professor in Computer Science at University of Montreal. His PhD research (Indiana University Computer Science, 2000) investigated the dynamics of model neurons in response to music-like rhythmic patterns. He went on to do work in computational music cognition, focusing on music generation with recurrent neural networks, meter perception, beat tracking and expressive music performance. More recently he has focused on large-scale machine learning approaches to music recommendation, including work in learning informative representations of music audio. At Google he has been working on the music recommender for Music Beta by Google. Douglas Eck doesn't like bios written in the third person, but since everyone else does it, he does it too.